arXiv:1402.2095v1 [physics.comp-ph] 10 Feb 2014

Evaluating functions of positive-definite matrices using colored noise thermostats

Marco Nava,[1] Michele Ceriotti,[2] Chaim Dryzun,[1] and Michele Parrinello[1]

[1] *Computational Science, Department of Chemistry and Applied Biosciences, ETH Zurich, USI-Campus, Via Giuseppe Buffi 13, C-6900 Lugano, Switzerland*

[2] *Laboratory of Computational Science and Modeling, École Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland*

(Dated: February 11, 2014)

Many applications in computational science require computing the elements of a function of a large matrix. A commonly used approach is based on the the evaluation of the eigenvalue decomposition, a task that, in general, involves a computing time that scales with the cube of the size of the matrix. We present here a method that can be used to evaluate the elements of a function of a positive-definite matrix with a scaling that is linear for sparse matrices and quadratic in the general case. This methodology is based on the properties of the dynamics of a multidimensional harmonic potential coupled with colored noise generalized Langevin equation (GLE) thermostats. This "$f-$thermostat" (FTH) approach allows us to calculate directly elements of functions of a positive-definite matrix by carefully tailoring the properties of the stochastic dynamics. We demonstrate the scaling and the accuracy of this approach for both dense and sparse problems and compare the results with other established methodologies.

INTRODUCTION

Eigenvalue equations are ubiquitous in science, as they play an essential role in quantum mechanics [1–3], statistical mechanics [4], biology [5], graph theory [6], computer science [7], pure mathematics [8] and more [9].

In many cases, one evaluates the eigenvalues as an intermediate step towards computing a function f of the input matrix $\mathbf{M}$. In fact, the standard procedure to compute matrix functions is to diagonalize the input matrix $\mathbf{M}$ by finding its right eigenvectors $\mathbf{u}_i$ and its eigenvalues ϵ_i. If $\mathbf{U}$ is the matrix that has $\mathbf{u}_i$ as columns and $\boldsymbol{\epsilon} = \mathrm{diag}\{\epsilon_i\}$ is a diagonal matrix with the eigenvalues on the diagonal, one can write

$$\mathbf{M} = \mathbf{U}\boldsymbol{\epsilon}\mathbf{U}^{-1}, \quad \text{and} \quad f(\mathbf{M}) = \mathbf{U}\mathrm{diag}\left\{f(\epsilon_i)\right\}\mathbf{U}^{-1}. \quad (1)$$

The most significant shortcoming of this approach is that conventional techniques for computing the eigenvalue decomposition of a matrix of size N have a computational cost that scales with the cube of the size of the matrix, which makes them very inefficient when used on large matrices [10].

In many applications, however, one is dealing with very sparse matrices, and needs to compute just a small subset of the elements of $f(\mathbf{M})$. In these cases, one should consider whether it is possible to obtain algorithms with a more favorable scaling. A typical example can be found in the field of electronic structure calculations, where the interaction potential between the particles can be obtained by evaluating the density matrix $\boldsymbol{\rho}$ – that at finite temperature is just the Fermi function of the Hamiltonian matrix $\mathbf{H}$ – and then computing $\mathrm{Tr}\,\boldsymbol{\rho}(\mathbf{H})\mathbf{H}$.

The efficiency of this kind of calculations for large systems can be improved greatly if one can evaluate $\boldsymbol{\rho}$ while avoiding the diagonalization of the Hamiltonian, using instead methods with a cost that scales linearly with the size of the system. Many different linearly scaling algorithms have been suggested [2, 3, 11–23], but they are still not used routinely, in part because of their complexity and their memory requirements, and in part because they only become advantageous for very large system sizes. In this paper we introduce a method based on colored-noise, generalized Langevin equation (GLE) thermostats [24, 25], that can be used to compute selected elements of the function of a symmetric, positive-definite sparse matrix with linear-scaling effort.

Our approach is a generalization of the ideas underlying the so-called δ-thermostat [26]. By considering the input matrix as the Hessian of an N-dimensional harmonic oscillator, one can generate an artificial, non-canonical GLE dynamics that selectively enforces different occupations onto the various normal modes. The key idea behind the f-thermostat is that the GLE thermostat can be tailored to generate an arbitrary frequency dependence of the fluctuations of position and momentum, and that with an appropriate choice of such dependence one can obtain the elements of a function of the input matrix as averages of the dynamical variables. We will start by describing the details of the methodology, and then present its application to different symmetric matrices and real functions, that will be used to discuss the advantages of the method, the statistic and systematic errors, and demonstrate the linear scaling of its computational cost for large, sparse matrices.

I. THE F-THERMOSTAT

A non-Markovian stochastic differential equation (SDE), based on the Langevin equation, for a particle with mass-scaled position q and mass-scaled momentum p, moving under the influence of a potential $V(q)$ can be

written as [27–30]:

$$\begin{aligned}\dot{q} &= p \\ \dot{p} &= -\frac{\partial V(q)}{\partial q} - \int K(t-s)\, p(s)\, ds + \zeta(t)\,.\end{aligned} \tag{2}$$

Here $K(t)$ is a memory kernel describing the friction, and $\zeta(t)$ is a noisy-force term characterized by Gaussian statistics and time correlation function $H(t) = \langle \zeta(t)\,\zeta(0)\rangle$. Treating analytically and integrating numerically the SDE in this form is inconvenient. Generating a sequence of random numbers with the prescribed correlation is possible but not straightforward [31]. Furthermore, storing and referring to the past momenta in order to introduce the friction memory kernel is inefficient from the point of view of memory requirements and computation time.

However, Eq. (2) can be represented [32–34] by a Markovian, Langevin equation in an extended momentum space containing a certain number of additional momenta $\{s_i\}$. One can then write [25, 35],

$$\begin{aligned}\dot{q} =& p \\ \begin{pmatrix}\dot{p}\\ \dot{\mathbf{s}}\end{pmatrix} =& \begin{pmatrix}-V'(q)\\ \mathbf{0}\end{pmatrix} - \begin{pmatrix}a_{pp} & \mathbf{a}_p^T\\ \bar{\mathbf{a}}_p & \mathbf{A}\end{pmatrix}\begin{pmatrix}p\\ \mathbf{s}\end{pmatrix} + \begin{pmatrix}b_{pp} & \mathbf{b}_p^T\\ \bar{\mathbf{b}}_p & \mathbf{B}\end{pmatrix}\begin{pmatrix}\boldsymbol{\xi}\end{pmatrix},\end{aligned} \tag{3}$$

where $\boldsymbol{\xi}$ is a vector of $n+1$ uncorrelated Gaussian random numbers satisfying $\langle \xi_i(t)\,\xi_j(t')\rangle = \delta_{ij}\delta(t-t')$ and $\langle \xi_i\rangle = 0$. By integrating out the additional fictitious degrees of freedom, the non-Markovian Langevin equation (2) is recovered describing the trajectory of the physical degrees of freedom (q,p), with memory kernels that correspond to a fairly general superimposition of complex exponentials [25, 35]. In this Markovian form, the dynamics is determined by the drift and diffusion matrices

$$\mathbf{A}_p = \begin{pmatrix}a_{pp} & \mathbf{a}_p^T\\ \bar{\mathbf{a}}_p & \mathbf{A}\end{pmatrix} \quad \text{and} \quad \mathbf{B}_p = \begin{pmatrix}b_{pp} & \mathbf{b}_p^T\\ \bar{\mathbf{b}}_p & \mathbf{B}\end{pmatrix} \tag{4}$$

respectively. It is often more convenient to characterize the dynamics in terms of the free-particle static covariance $\mathbf{C}_p = \langle (p,\mathbf{s})(p,\mathbf{s})^T\rangle$, i.e. the matrix that contains the stationary correlations between p and $\mathbf{s}$ in the absence of an external potential. $\mathbf{C}_p$ is related to $\mathbf{B}_p$ and $\mathbf{A}_p$ through

$$\mathbf{A}_p\mathbf{C}_p + \mathbf{C}_p\mathbf{A}_P^T = \mathbf{B}_p\mathbf{B}_p^T, \tag{5}$$

and can be used to express a number of physical constraints on the dynamics and to derive a relatively straightforward integrator for Eqs. (3), based on a modified velocity-Verlet algorithm [25, 35].

What makes the formalism of Eq. (3) particularly useful is that the sampling and dynamical properties of the GLE trajectory can be predicted analytically in the case of a harmonic potential $V(q) = \omega^2 q^2/2$. This made it possible to obtain within the same formalism a number of useful effects, ranging from efficient canonical sampling [24, 36] to the stabilization of multiple time step dynamics [37] and the modeling of nuclear quantum effects [38–40].

Within this framework, one can obtain generalized Langevin dynamics that do not fulfill the classical fluctuation-dissipation theorem, and therefore do not sample the canonical ensemble at a given temperature. In the harmonic limit, this stochastic dynamics can be understood in terms of a steady-state process in which normal modes of different frequencies reach a stationary distribution in which position and momentum have Gaussian statistics, with frequency-dependent fluctuations. Manipulating the behavior of $\langle q^2\rangle(\omega)$ and $\langle p^2\rangle(\omega)$ is the basic idea behind the "quantum thermostat"[38], the δ-thermostat [26], and the method we introduce in this paper. Indeed, the f-thermostat can be viewed as a generalization of the δ-thermostat, which uses non-canonical sampling to thermalize at a higher effective temperature only the characteristic frequencies of the system that lie in a prescribed, narrow frequency range. The other frequencies equilibrate at a much lower effective temperature. This is done by enforcing a delta-like frequency dependence of the fluctuations through a GLE based thermostat, and can be used to estimate the density of states of a positive-definite matrix [26]. The idea behind the present method is that by a careful choice of the frequency dependence of the fluctuations one can directly evaluate the elements of $f(\mathbf{M})$.

To see how this can be achieved, let us consider a symmetric, positive-definite matrix, $\mathbf{M}$ of size N. This matrix can be used as the Hessian of a multi-dimensional harmonic potential:

$$V(\mathbf{q}) = \frac{1}{2}\mathbf{q}^T\mathbf{M}\mathbf{q}. \tag{6}$$

One could then study the dynamics of a particle with unit mass subject to this harmonic potential. The resulting problem is best studied in terms of normal mode coordinates: if the k-th eigenvalue of $\mathbf{M}$ is ϵ_k, and the associated eigenvector $\mathbf{u}_k$, one can write the displacement along the normal mode as $\tilde{q}_k = \mathbf{u}_k^T\mathbf{q}$ and the corresponding momentum as $\tilde{p}_k = \mathbf{u}_k^T\mathbf{p}$. If the dynamics is microcanonical, this normal mode will oscillate indefinitely with a frequency $\omega_k = \sqrt{\epsilon_k}$. If instead the system is coupled to a heat bath at temperature T – for instance by using a Langevin equation – the fluctuations of displacement and momentum are bound to be $\langle \tilde{q}_k^2\rangle = T/\omega_k^2$ and $\langle \tilde{p}_k^2\rangle = T$, taking for simplicity $k_B = 1$.

Now let us consider a GLE thermostat whose parameters have been fitted in such a way that, over the range of frequencies that corresponds to the spectral range of $\mathbf{M}$, a harmonic oscillator of frequency ω will have momentum fluctuations $\langle p^2\rangle(\omega) = T^\star(\omega)$, where we introduced the effective temperature $T^\star$ as a shorthand for the average momentum fluctuations. Let us assume that the GLE parameters have been chosen so that the momentum fluctuations are related to the desired matrix function f as

$T^\star(\omega) = f\left(\omega^2\right)$. If we then apply an independent thermostat of this kind to each of the N "Cartesian" degrees of freedom of the artificial dynamics generated by the potential (6), each normal mode will respond as if the GLEs were applied in the normal modes representation. This is a consequence of the invariance of the combined GLE dynamics with respect to a unitary transformation of the coordinates, which in turn is guaranteed by the Gaussian nature of the noisy force and by the linear nature of the stochastic differential equations [35]. The individual normal modes described by the eigenvalues of $\mathbf{M}$ will then respond to the GLE dynamics and reach a steady state with fluctuations that are consistent with the thermostat specifications; namely $\langle \tilde{p}_k \tilde{p}_{k'} \rangle = \delta_{kk'} T^\star(\omega_k) = \delta_{kk'} f(\epsilon_k)$. Since the eigenvectors u_k form an orthonormal basis and $p_i = \sum_k \tilde{p}_k U_{ik}$, one can write

$$\begin{aligned} \langle p_i p_j \rangle &= \sum_{kk'} U_{ik} U_{jk'} \langle \tilde{p}_k \tilde{p}_{k'} \rangle = \\ &= \sum_k U_{ik} f(\epsilon_k) U_{jk} = f(\mathbf{M})_{ij}. \end{aligned} \tag{7}$$

In particular, $\mathrm{Tr}\, f(\mathbf{M})$ can be obtained very easily from the kinetic energy of the GLE-thermostatted harmonic dynamics: $\mathrm{Tr}\, f(\mathbf{M}) = \sum_i \langle p_i^2 \rangle = \langle \mathbf{p}^T \mathbf{p} \rangle$.

Provided that $\mathbf{M}$ is sparse, the technique described above can be applied with an effort that scales linearly with the size of the matrix, because computing the forces which correspond to the potential (6) involves a matrix-vector product that can be evaluated in $\mathcal{O}(N)$. Averages that correspond to individual elements of the momentum covariance matrix can be evaluated with an effort independent of the size of $\mathbf{M}$, so the method remains linear scaling as long as one needs to compute a $\mathcal{O}(N)$ subset of the elements of $f(\mathbf{M})$. As we will discuss later, the relative statistical error on the results is also independent of N. The procedure we have highlighted here allows to obtain elements of $f(\mathbf{M})$ for large, sparse, positive-definite matrices without performing explicitly an eigenvalue decomposition, since both the energy (6) and the momentum covariance (7) can be computed in the "Cartesian" representation. In a similar way one can also obtain traces of the form $\mathrm{Tr}\, \mathbf{O} f(\mathbf{M})$, which often appear when computing observables in electronic structure calculations and when evaluating thermal averages. In fact, it is easy to see that $\mathrm{Tr}\, \mathbf{O} f(\mathbf{M}) = \left\langle \mathrm{Tr}\, \mathbf{O} \mathbf{p} \mathbf{p}^T \right\rangle = \left\langle \mathbf{p}^T \mathbf{O} \mathbf{p} \right\rangle$.

While the idea behind this approach is relatively simple, it is also true that designing a thermostat that enforces the desired $\left\langle p^2 \right\rangle(\omega)$ dependence is not trivial [25, 35]. Many conflicting goals must be fulfilled besides the prescribed ω-dependent momentum fluctuations, e.g. the stability of the dynamics and the efficiency of the sampling. The fitting strategy we use is based on that described in Refs. [25, 35], but we found advantageous to start with a simplified parametrization of the $\mathbf{A}_p$ and $\mathbf{C}_p$ matrices that corresponds to a combination of δ-like memory kernels [26, 35]. The resulting matrices can then be fine-tuned using the more general parametrization described in Ref. [25]. In practice, obtaining an effective set of parameters requires some experience and a few hours of work. However, as discussed in previous publications, once a pair of $\mathbf{A}_p$ and $\mathbf{C}_p$ matrices has been obtained for a given application, it is possible to re-use them in many similar circumstances, by modifying the range of frequencies they work on just by scaling their value.

In the case of the f-thermostat, scaling the GLE parameters can also be used to extend the range of functions that can be computed. In fact, for the artificial dynamics to yield a well-defined stationary distribution $\mathbf{M}$ must be positive-definite, and f must be positive-valued, so that the method is restricted to the evaluation of $f : D \subseteq \mathbb{R}^+ \to C \subseteq \mathbb{R}^+$. Shifting and scaling $\mathbf{A}_p$, $\mathbf{C}_p$ and $\mathbf{M}$ can be used to change the spectral range of the input matrix to make it positive-definite, while still obtaining the desired matrix function.

Let us consider the scaling of the matrices in Eq. (5) and (4), $\mathbf{A}'_p = \alpha \mathbf{A}_p$ and $\mathbf{C}'_p = \gamma \mathbf{C}_p$, and the transformation of $\mathbf{M}$ into $\mathbf{M}' = r\mathbf{M} + m\mathbf{I}$ – where $\mathbf{I}$ is the $N \times N$ identity matrix. If $\mathbf{M}$ has eigenvalues within the range $[\epsilon_{\min}, \epsilon_{\max}]$, then one can obtain a transformed matrix with range $[\epsilon'_{\min}, \epsilon'_{\max}]$ by choosing $r = \left(\epsilon'_{\min} - \epsilon'_{\max}\right) / \left(\epsilon_{\min} - \epsilon_{\max}\right)$ and $m = \left(\epsilon_{\max}\epsilon'_{\min} - \epsilon_{\min}\epsilon'_{\max}\right) / \left(\epsilon_{\min} - \epsilon_{\max}\right)$. In the most general case, when all the transformations are applied simultaneously, the f-thermostat procedure would yield $\left\langle \mathbf{p}\mathbf{p}^T \right\rangle = f'(\mathbf{M}') = \gamma f\left((r\mathbf{M} + m\mathbf{I})/\alpha^2\right)$. Depending on the function f it might still be possible to perform the inverse transformation and recover the original target. For instance, if $f(x) \equiv \ln(x)$, then $f\left(\mathbf{M}\right) = f\left(r\mathbf{M}\right) - \mathbf{I} \cdot \ln r$, and if $f(x) \equiv \exp(x)$, then $f\left(\mathbf{M}\right) = f\left(\mathbf{M} + m\mathbf{I}\right) e^{-m}$,

Let us stress from the start that this does not aim to be a very accurate method. A first source of error is the fact that the fit will not in general be capable of enforcing perfectly the desired behavior for $\left\langle p^2 \right\rangle(\omega)$. More flexibility in the fit – and a tighter compliance to the desired constraints – can be achieved to an extent by increasing the number n of additional degrees of freedom, at the cost however of larger memory and computational requirements. Typically, we find that a choice between 4 and 10 additional momenta $\{s_i\}$ provides sufficient flexibility to obtain a maximum relative error of one percent or less over a range of frequencies spanning several orders of magnitude. Other errors can originate from the parameters of the simulation and can be in principle be made arbitrarily small. For instance, finite time-step errors can be minimized by choosing small enough time steps; when using a velocity Verlet integrator [25], one could also perform the fit for a specific integration time-step, modifying the target fluctuations vs. frequency curve in such a way that finite time-step errors are accounted for and exactly corrected. The most serious limitation of this method, however lies in its stochastic nature. The elements of $f(\mathbf{M})$ are obtained from an average of dynamical variables, and the average will be affected by a statistical error. This component of the error decreases slowly with simulation time, specifically with the square root of the number of uncorrelated samples. The method we propose

here is best suited for the (many) applications in which a non-negligible stochastic error can be tolerated and, perhaps, even useful [17, 18], or to initialize other approaches that converge faster but require an initial guess for $f(\mathbf{M})$.

II. RESULTS AND BENCHMARKS

As discussed, the f-thermostat relies on the accuracy of the fit of the frequency-dependent fluctuations to the desired target; however the quality of the results as well as the performance of the methodology depend also on other constraints that we will discuss here briefly. In summary, a successful set of colored-noise parameters should guarantee: (1) a tight agreement between $T^\star(\omega)$ and $f(\omega^2)$; (2) stability of the dynamics with respect to the integration time step; (3) statistical efficiency.

The first requirement is rather obvious since the discrepancy in the fit of $T^\star(\omega)$ will cause systematic errors in the evaluated matrix function. It is not hard to obtain small relative errors over a large range of frequencies for slowly-varying functions. However, the frequency-dependence of $T^\star(\omega)$ for simple forms of the GLE matrices have typically an algebraic $\propto 1/\omega^2$ decay for high values of ω. This is particularly problematic when one wants to obtain an exponentially-decaying $T^\star(\omega)$. In many applications it is sufficient that the small $f(\epsilon)$ are negligible with respect to the larger ones. Therefore, when fitting exponential tails, we have decided to choose a small threshold δ, and just require that for large ω the effective temperature of normal modes is smaller than $\delta \max\left[T^\star(\omega)\right]$ – rather than it matches the exponential decay.

The second requirement has an impact on the efficiency of the simulation. The time discretization dt of the dynamics in the absence of noise is determined by the largest eigenvalue of $\mathbf{M}$, i.e. $dt \ll 1/\sqrt{\epsilon_{\max}}$. However, if the noise contains strong components with $\omega > \sqrt{\epsilon_{\max}}$, it may become necessary to reduce the time step further, in order to integrate properly the stochastic dynamics and obtain the fluctuations that are predicted analytically for an infinitesimal time step. In practice, this constraint corresponds to making sure that the Fourier transform of the memory kernels $K(\omega)$ and $H(\omega)$ do not exhibit prominent features for $\omega > \sqrt{\epsilon_{\max}}$, and that $\lim_{\omega\to\infty} K(\omega) < \sqrt{\epsilon_{\max}}$. These requirements can be implemented by including additional penalty terms in the fitness function, as discussed in Ref. [25].

The third requirement determines how quickly the stochastic dynamics generates uncorrelated samples of positions and momenta, and is therefore essential to speed up the convergence of averages with time. Since the method requires the calculation of $\langle p_i p_j \rangle$, it is important to minimize the correlation time of p^2, $\tau_{p^2}(\omega)$, over the range of frequencies that is relevant for the problem at hand. If $t_{\rm tot}$ is the total simulation time, the statistical error on the value of $\langle p^2 \rangle$ for a normal mode

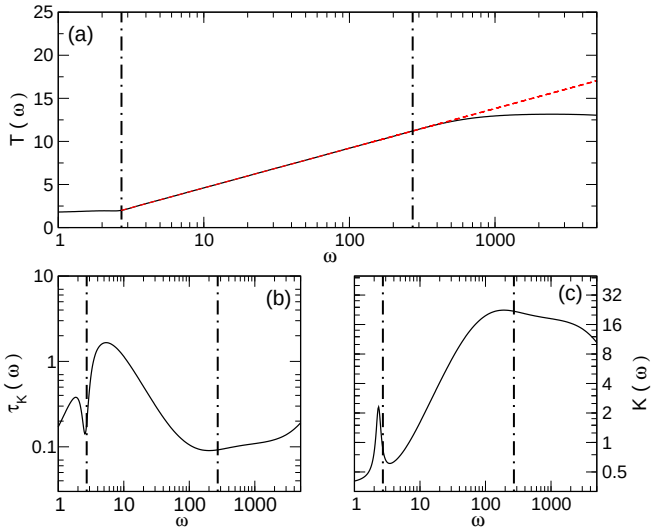

Figure 1. (Color online) The fit of the fluctuations versus frequency dependence of the thermostat along with the constraints required for an efficient thermostatting of the system. (a) $\langle p^2 \rangle$ dependency and (Dashed line) the fit target, the fit interval is marked by the vertical dotted lines; (b) Autocorrelation time of p^2 as a function of frequency; (c) Effective friction – the Fourier transform of the memory kernel $K(\omega)$. Arbitrary units have been used.

of frequency ω will be of the order of $\sqrt{\tau_{p^2}(\omega)/t_{\rm tot}}$ – so minimizing $\tau_{p^2}(\omega)$ allows one to obtain smaller statistical uncertainty for a given length of the simulation. A small autocorrelation time is also necessary to make sure that the final result will be independent on the initial conditions of the dynamics – which could in principle introduce a systematic bias in the estimate of $f(\mathbf{M})$.

A remarkable feature of the FTH method is the very low memory consumption. If $\mathbf{M}$ is sparse, the memory footprint grows linearly with the size of the matrix: the storage required consists of roughly $M + (3+S)N + 2A$ real numbers. The first term is required for the storage of matrix $\mathbf{M}$ in sparse format – corresponding to the number M of non-zero entries; the second term is the memory required for the modified Verlet algorithm, where N is the size of $\mathbf{M}$ and S is the number of additional momenta, and the last term is the storage required to collect and block-average an arbitrary number A of elements of $f(\mathbf{M})$ that one wishes to compute. The memory required by the FTH method in the case of a dense matrix is also relatively small and grows quadratically with the size of $\mathbf{M}$: one would need $N(N+1)/2$ elements to store $\mathbf{M}$, another $N(N+1)$ for the block-averaging of the entire $f(\mathbf{M})$ and a linear-scaling $(3+S)N$ for the propagation of the GLE dynamics.

Let us now present a few examples, where we apply the f-thermostat to compute various matrix functions for matrices of different form.

Matrix logarithm. We chose as a first example the problem of evaluating the logarithm of a matrix. Fig. 1 shows the $\langle p^2 \rangle(\omega)$ curve for an optimized set of parame-

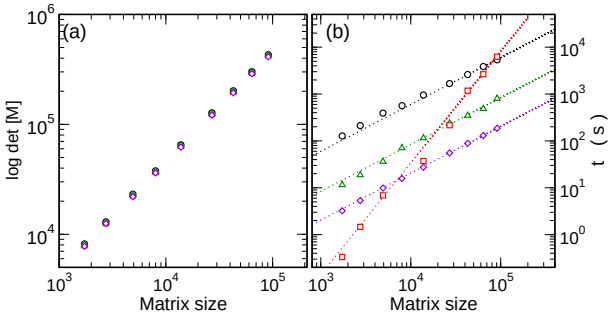

Figure 2. (Color online) Panel a) logarithm of the determinant of a sparse matrix $\mathbf{M}$ obtained from the trace of the logarithm of $\mathbf{M}$ with a Sparse Cholesky decomposition (Squares) and the f-thermostat approach with a 0.1% (Circles), 1% (Triangles) and 2.5% (Diamonds) uncertainty on the results. Panel b) The computing time required to evaluate the quantities displayed in panel (a). Symbols have the same meaning as those in panel (a). Dotted lines are power-law fits of the computing time.

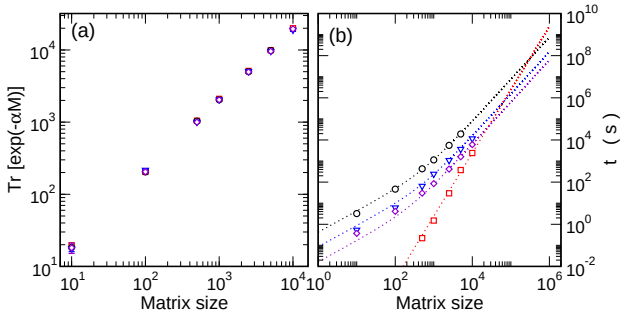

Figure 3. (Color online) Panel a) (Squares) Tr[exp($-\alpha\mathbf{M}$)] obtained with a diagonalization approach [42]. Other symbols: the same obtained with the f-thermostat approach with uncertainties of 0.2% (Circles), 0.5% (Triangles) and 1% (Diamonds) on the results. Panel b) Time scaling for the computations in panel (a). Symbols have the same meaning of those in panel (a). Dotted lines are fit to computing time data.

ters, following the desired $\ln\omega^2$ behavior for frequencies within the interval $[e, 100e]$. The figure also shows the sampling efficiency and the friction kernel, whose importance has been discussed in the previous paragraph.

In Fig. 2 we show the values obtained for the evaluation of the trace of the logarithm of matrices $\mathbf{M}_1$ of different size, using the colored noise parameters described in Fig. 1, together with the wall-clock time needed to converge the results with different levels of accuracy. The results and timings are compared with those obtained with a supernodal sparse Cholesky decomposition[41] of $\mathbf{M}_1$, an highly optimized tool for the computation of determinants of symmetric positive definite matrices that, through the relation $\ln\det(\mathbf{M}) = \mathrm{Tr}\ln\mathbf{M}$, allows for a check of the results. Here $\mathbf{M}_1$ matrices have the form of the tight binding Hamiltonian of a three dimensional cubic lattice with periodic boundary conditions with up to the sixth nearest neighbors random hopping terms and a constant on-site interaction that keeps the matrix positive definite. Details on how $\mathbf{M}_1$ matrices have been defined are reported in the Appendix.

Matrices of this form are frequently encountered in physics and it is known that, in general, it is not possible to transform them into banded matrices that can be manipulated with linear scaling effort by sparse linear algebra methods. In fact, the time required for the supernodal Cholesky decomposition (SC) scales superquadratically with matrix size, as $\mathcal{O}(N^{2.4})$. As expected, the FTH method scales linearly with system size. For small matrices the SC method is faster – SC is a highly optimized method specifically aimed at computing a matrix determinant, so a direct comparison with the more general-purpose FTH is somewhat unfair – but its worse system-size scaling will inevitably lead to the FTH method to prevail for large enough N. As it is apparent from Fig. 1, the break-even point depends dramatically on the desired level of statistical accuracy. For a relative error of the order of 0.1%, the break-even point is at $N \sim 10^5$.

Matrix exponential. The FTH methodology can also be used in the case of dense matrices, where its computational cost grows as $\mathcal{O}(N^2)$. Fig. 3 shows a comparison between the FTH method and a diagonalization method [42] for the computation of the exponential of a dense symmetric random matrix $\mathbf{M}_2$ with eigenvalues in the interval $(a, b) \approx (-300, 300)$; namely $\exp(-\alpha\mathbf{M}_2)$ with $\alpha = 1/100$ (see the Appendix for more details on the form of $\mathbf{M}_2$).

The matrix exponential is a typical operation encountered in many fields of physics – for instance when computing thermal density matrices $\exp(-\beta\hat{H})$. For this operation, the conventional diagonalization approach scales with the cube of the matrix size, so again for large enough matrices the FTH method will prevail. Note that for small N the cost of the FTH method grows linearly with matrix size, as it is dominated by the (linear scaling) cost of propagating the stochastic dynamics and not by the evaluation of the force.

Comparing the timings reported in Fig. 1 and Fig. 3 in the size range in which the propagation of the dynamics dominates the computational cost of the FTH, one can see that for a given accuracy evaluating $\exp(\mathbf{M})$ is much more demanding than evaluating $\ln(\mathbf{M})$. This is a consequence of the fact that the autocorrelation time τ_{p^2} we could obtain for the exponential fit was longer than in the case of the logarithm – so a longer artificial dynamics is necessary to obtain the same statistical uncertainty. This in turn is a consequence of the difficulty of obtaining the sharp decay which is characteristic of the exponential function, as explained above. A better efficiency could be obtained by using a larger threshold for the exponential tail (here we tried to fit the exponential decay down to $e^{-x} \approx 10^{-4}$), or by targeting specifically the range of frequencies that is relevant to the application at hand.

Fermi function. The examples reported this far had the purpose of benchmarking the methodology and did

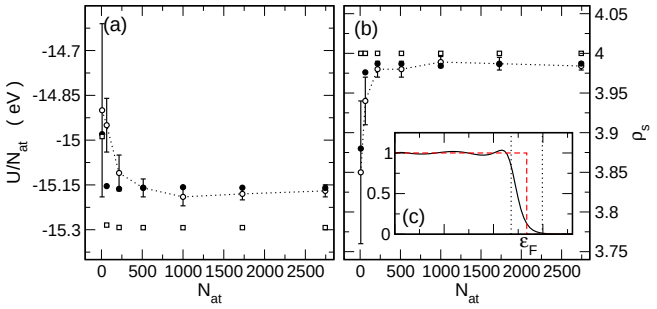

Figure 4. (a) Internal electron energy per atom and (b) average electron number per atom of a diamond bulk carbon lattice with lattice constant $l = 3.57$ Å and temperature $T \to 0$ K. The Fermi energy level has been placed in the middle of the energy gap so that $N_e/N = 4$. Open circles refer to the results obtained with the f-thermostat method, closed circles are the same data obtained by exact diagonalization but computing the approximate fit to the Fermi function (i.e. they correspond to the result predicted for the GLE dynamics in the absence of time step and statistical error). Open squares are the exact results computed from the eigenvalues using a step function. (c) (full line) The approximate Fermi function obtained from the fluctuation profile $\langle p^2 \rangle(\omega)$ of the f-thermostat we used for this example; the dashed (red) line marks the zero-temperature Fermi function and the vertical dotted lines show the boundaries of the energy gap.

not have any direct physical relevance. We show now an application of the FTH method to a concrete physical problem – namely, the evaluation of the electrons per site ρ_s and the energy U/N_{at} for a cubic supercell containing N_{at} carbon atoms in a diamond lattice, with lattice constant $l = 3.57$ Å, and at zero temperature, described by a tight binding Hamiltonian[43] with only first and second nearest neighbors interactions.

Given the tight binding Hamiltonian matrix H that describes the valence electrons of carbon, U/N_{at} and ρ_s can be computed with occupations given by the Fermi function $f_F(\epsilon)$,

$$f_F(\epsilon) = \frac{1}{e^{\beta(\epsilon-\epsilon_F)}+1} \tag{8}$$

$$\rho_s = 2\,\mathrm{Tr}\left[f_F(H)\right]/N_e \tag{9}$$

$$U/N_{at} = 2\,\mathrm{Tr}\left[H \cdot f_F(H)\right]/N_{at} \tag{10}$$

where the factor of 2 takes into account the spin degeneracy of the electron, ϵ_F is the Fermi energy of the system and N_e is the number of electrons,

In the low-temperature limit, $f_F(\epsilon) \to \theta(\epsilon_F - \epsilon)$ (a Heaviside theta function), and would be very difficult to reproduce with a f-thermostat, because one would have to fit $T^\star(\omega)$ to a sharp, discontinuous target. In the case of an insulator, however, it is not necessary to reproduce precisely f_F, since ϵ_F falls in the energy gap and thus a smooth transition of $T^\star(\omega)$ between 0 and 1 does not affect results, as long as it is confined to a neighborhood of ϵ_F entirely contained in the gap. This is the case shown in Fig. 4c; the energy gap for a diamond lattice of carbon atoms, in fact, is roughly 6.5 eV and for the computation of properties that do not involve a shift in the Fermi level, ϵ_F can be safely placed in the middle of the gap; the frequency-dependent fluctuations have been fitted with a maximum error outside the gap of 3%, the most relevant discrepancy being the oscillation shortly before the energy gap. In the case of the carbon Hamiltonian, this introduces a systematic error of about 1% on the values of ρ_s and U, that is clearly visible in Fig. 4b. Let us stress that if one needed to perform a self-consistency procedure, or to modify ϵ_F to enforce constant electron number when considering different structures as it would be the case when optimizing the geometry or running a molecular dynamics simulation, it would be possible to just scale the $\mathbf{A}_p$ and/or $\mathbf{M}$ matrices without having to perform a new fit for each configuration.

All data in Fig. 4a and 4b has been obtained with the same parameters for the Verlet dynamics. The relative statistical error is stable with respect to the size of the matrix, which is clearly a desirable feature of a method aimed at large sparse matrix calculations. Note in passing that the f-thermostat does not require an initial guess for $f(\mathbf{M})$, contrary to other methods to compute matrix functions in general, and the density matrix in particular. For instance, a widely used approach to compute the zero-temperature Fermi function is to perform Newton–Schulz iterations [44] that converge to the sign function. This procedure has a convergent behavior only if the starting point is already a good estimate of $f(\mathbf{M})$. In this and similar cases the f-thermostat can provide a sufficiently accurate estimate of $f(\mathbf{M})$ to be used as starting point of the iteration.

III. FINAL DISCUSSION

In this work we have used the properties of a generalized Langevin dynamics coupled to a multi dimensional harmonic oscillator in order to compute the elements of a function of a positive-definite matrix. We have verified that this approach entails a linear-scaling computational cost when the input matrix is sparse, and a quadratic cost in the case of a dense input matrix. It is therefore an efficient method when one needs to compute a function of a very large matrix.

Tuning the parameters of the GLE thermostat is a very delicate step. The frequency-dependent effective temperature $T^\star(\omega)$ has to be compatible with the target function in the relevant frequency range; at the same time, the strength of the noise and the autocorrelation time of p^2 should be minimized to guarantee stable dynamics and ergodic sampling. Discrepancies in the fit introduce systematic errors, whereas the other constraints affect also the performances of the computation. In particular, a small autocorrelation time reduces the statistical error for a given length of the simulation.

The stochastic nature of the FTH method implies that the result will be affected by a slowly-decaying statistical

error – the simulation time must be multiplied by four just to half the statistical uncertainty. There is a trade-off between accuracy and computational cost, that in practice will make this idea applicable only in cases where a relatively large error, of about 1%, can be tolerated – for instance when the results are to be used as the forces in a Langevin dynamics, or as the initial guess for a more accurate iterative method. This said, the FTH method offers three main advantages: on top of the favorable scaling of computational time with the size of the input matrix, one should also consider the very small memory footprint, and the ease of parallelization – not only the matrix-vector multiplications can be parallelized, but one could perform independent, trivially-parallel simulations to reduce the statistical error.

Appendix A: Details on the form of the example matrices

Here we report the prescription for generating the $\mathbf{M}_1$ and $\mathbf{M}_2$ matrices that we used in our examples. The parameters of the f-thermostat that we used are provided in the Supporting Information [**?**], together with a simple program that performs the artificial GLE dynamics and computes the elements of $f(\mathbf{M})$.

$\mathbf{M}_1$ *matrices (logarithm example).* The on-site terms have been set to $U = 100$ (arbitrary units) and the hopping terms are random numbers between 0 and 2. Having fixed a number n of sites per directions of the lattice, the diagonal elements (i, i) correspond to the on-site terms, whereas the off-diagonal elements (i, j) of the matrix correspond to the hopping terms between the elements $(i_x, i_y, i_z) = (i/n^2, (i\%n^2)/n, (i\%n^2)\%n)$ and $(j_x, j_y, j_z) = (j/n^2, (j\%n^2)/n, (j\%n^2)\%n)$ of the lattice, where the divisions are considered in the integer number field and % is intended as the division remainder operation. This naturally defines the distance of two elements of the matrix, $m_1 = (i_0, j_0)$, $m_2 = (i_1, j_1)$ as the distance in periodic boundary conditions of the corresponding sites on the lattice; from this distance one can then determine any k-th neighbor of a lattice site.

$\mathbf{M}_2$ *matrices (exponential).* The random matrices have been obtained in a two step process; consider a $m \times m$ square matrix: first, the diagonal elements $D(1)$,$D(2)$,...,$D(m)$, have been determined with a uniform random number between $a + (b-a)/4$ and $b - (b-a)/4$ for each diagonal element. Then, each off-diagonal element is determined in the lower triangular part of the matrix; the eigenvalues interval is controlled with the Gerschgorin circle theorem in the following way: if the element (i, j) is x, then it will contribute to the Gerschgorin radii $R(i)$ and $R(j)$ of both the i and j row due to the symmetry of the matrix; more precisely, adding an off-diagonal element x at position (i, j) will increase both $R(i)$ and $R(j)$ by a positive quantity $|x|$. Fixing a row i, each element $j_0, j_1, ..., j_{i-1}$ is a random number uniformly distributed between $-r/(i-j)$ and $r/(i-j)$, where $r = \min[D(i) - R(i) - a, b - D(i) - R(i), D(j) - R(j) - a, b - D(j) - R(j)]$; the division by $i - j$ ensures an almost complete filling of the matrix elements. This procedure guarantees that the eigenvalues of the resulting matrix lie in the interval (a, b).

[1] S. Bernardson, P. McCarty, and C. Thron, Computer physics communications **78**, 256 (1994).
[2] W. Kohn, Phys. Rev. Lett. **76**, 3168 (1996).
[3] S. Goedecker, Rev. Mod. Phys. **71**, 1085 (1999).
[4] A. Alavi and D. Frenkel, J. Chem. Phys. **97**, 9249 (1992).
[5] M. Porto, U. Bastolla, H. E. Roman, and M. Vendruscolo, Physical review letters **92**, 218101 (2004).
[6] B. Bollobás, *Modern graph theory*, Vol. 184 (Springer, 1998).
[7] K. Bryan and T. Leise, Siam Review **48**, 569 (2006).
[8] Z. Bai, G. Fahey, and G. Golub, Journal of Computational and Applied Mathematics **74**, 71 (1996).
[9] P. Bonacich, American journal of sociology , 1170 (1987).
[10] G. H. Golub and C. F. Van Loan, *Matrix computations*, Vol. 3 (JHU Press, 2012).
[11] W. Yang, Phys. Rev. Lett. **66**, 1438 (1991).
[12] G. Galli and M. Parrinello, Phys. Rev. Lett. **69**, 3547 (1992).
[13] X. P. Li, R. W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10891 (1993).
[14] S. Goedecker and L. Colombo, Phys. Rev. Lett. **73**, 122 (1994).
[15] A. H. R. Palser and D. E. Manolopoulos, Phys. Rev. B **58**, 12704 (1998).
[16] D. R. Bowler and M. J. Gillan, Comp. Phys. Comm. **120**, 95 (1999).
[17] F. R. Krajewski and M. Parrinello, Phys. Rev. B **71**, 233105 (2005).
[18] F. Krajewski and M. Parrinello, Phys. Rev. B **73**, 041105 (2006).
[19] F. R. Krajewski and M. Parrinello, Phys. Rev. B **74**, 125107 (2006).
[20] F. R. Krajewski and M. Parrinello, Phys. Rev. B **75**, 235108 (2007).
[21] M. Ceriotti, T. D. Kühne, and M. Parrinello, J. Chem. Phys. **129**, 24707 (2008).
[22] M. Ceriotti, T. D. Kühne, M. Parrinello, and T. D. Kiihne, AIP Conference Proceedings **1148**, 658 (2009).
[23] L. Lin, J. Lu, R. Car, and E. Weinan, Phys. Rev. B **79**, 115133 (2009).
[24] M. Ceriotti, G. Bussi, and M. Parrinello, Phys. Rev. Lett. **102**, 020601 (2009).
[25] M. Ceriotti, G. Bussi, and M. Parrinello, J. Chem. Theory Comput. **6**, 1170 (2010).
[26] M. Ceriotti and M. Parrinello, Procedia Computer Science **1**, 1607 (2010).
[27] B. J. Berne and D. Forster, Annual Review of Physical Chemistry **22**, 563 (1971).

[28] R. F. Fox, J. Stat. Phys. **46**, 1145 (1987).
[29] Y. Abe, S. Ayik, P. G. Reinhard, and E. Suraud, Physics Reports **275**, 49 (1996).
[30] J. uczka, Chaos **15**, 26107 (2005).
[31] A. A. Maradudin, T. Michel, A. R. McGurn, and E. R. Méndez, Ann. Phys. **203**, 255 (1990).
[32] R. Zwanzig, J. Stat. Phys. **9**, 215 (1973).
[33] R. Zwanzig, *Nonequilibrium statistical mechanics* (Oxford University Press, New York, 2001).
[34] R. Zwanzig, Phys. Rev. **124**, 983 (1961).
[35] M. Ceriotti, *A novel framework for enhanced molecular dynamics based on the generalized Langevin equation*, Ph.D. thesis, ETH Zürich (2010).
[36] M. Ceriotti, M. Parrinello, T. E. Markland, and D. E. Manolopoulos, J. Chem. Phys. **133**, 124104 (2010).
[37] J. A. Morrone, T. E. Markland, M. Ceriotti, and B. J. Berne, J. Chem. Phys. **134**, 14103 (2011).
[38] M. Ceriotti, G. Bussi, and M. Parrinello, Phys. Rev. Lett. **103**, 30603 (2009).
[39] M. Ceriotti, D. E. Manolopoulos, and M. Parrinello, J. Chem. Phys. **134**, 84104 (2011).
[40] M. Ceriotti and D. E. Manolopoulos, Phys. Rev. Lett. **109**, 100604 (2012).
[41] Y. Chen, T. A. Davis, W. W. Hager, and S. Rajamanickam, ACM Transactions on Mathematical Software (TOMS) **35**, 22 (2008).
[42] G. Guennebaud, B. Jacob, *et al.*, "Eigen v3," (2012).
[43] C. Xu, C. Wang, C. Chan, and K. Ho, Journal of Physics: Condensed Matter **4**, 6047 (1992).
[44] N. J. Higham, *Functions of matrices: theory and computation* (Siam, 2008).
[45] See Supplemental Material at [] for the used thermostat fit parameters and a basic, ready-to-use, implementation of the f-thermostat.